

KINETIC CLUSTERING OF POINTS ON THE LINE

CRISTINA G. FERNANDES, MARCIO T.I. OSHIRO

Instituto de Matemática e Estatística
Universidade de São Paulo, Brazil

ABSTRACT. The problem of clustering a set of points moving on the line consists of the following: given positive integers n and k , the initial position and the velocity of n points, find an optimal k -clustering of the points. We consider two classical quality measures for the clustering: minimizing the sum of the clusters diameters and minimizing the maximum diameter of a cluster. For the former, we present polynomial-time algorithms under some assumptions and, for the latter, a $(2.71 + \varepsilon)$ -approximation.

1. INTRODUCTION

Clustering refers to a well-known class of problems whose goal is to partition a set so that “similar” elements are placed in the same subset of the partition. The notion of similarity and the format of the partition depend on the application. In this work, we study two clustering problems in a kinetic context, where points move continuously.

Atallah [Ata85] proposed a model for the points movement where the points are in a d -dimensional space and each coordinate of each point is given by a polynomial on the time variable. Using his model, Har-Peled [HP04] showed how to apply a clustering algorithm for the static setting to find a competitive clustering of the moving points. His objective was to find k centers that cover all the points with minimum radius. When the polynomials describing the points movement have degree at most μ , his algorithm relaxes the restriction on the number of clusters, allowing at most $k^{\mu+1}$ clusters, in order to achieve a constant approximation ratio with respect to the optimal radius of a k -clustering at any time.

Another model, called KDS (kinetic data structure), is presented by Basch, Guibas, and Hershberger [BGH99]. In this model, there is no need to know the full description of the points movement, which can be updated online. Using this model, Gao, Guibas, Hershberger, Zhang, and Zhu [GGH⁺03] proposed a randomized constant approximation to maintain, as the points move, a clustering minimizing the number of discrete centers needed to cover all points within a fixed radius.

Lee, Han, and Whang [LHW07] presented a framework for clustering of trajectories, defined as a sequence of points in a multi-dimensional space describing the movement of

Date: Tuesday 15th December, 2015, 02:24.

This research has been partially supported by Capes, CNPq Projects (Proc. 308523/2012-1), Fapesp Project (Proc. 2013/03447-6) and MaCLinC Project of Numec/USP, Brazil.

Email: {cris|oshiro}@ime.usp.br.

an object. They reduced each trajectory to a set of contiguous line segments and used heuristics to group similar resulting trajectories.

Our work is closer to the one of Har-Peleg, as we look for a static clustering of moving points, instead of a way to keep an (almost) optimal k -clustering all the time. We consider a restricted case, with points located in \mathbb{R} , linear movements ($\mu = 1$), and two classical quality measures for the clustering: minimizing the sum of the clusters diameters and minimizing the maximum diameter of a cluster. For the former, we present a polynomial-time algorithm under some assumptions and, for the latter, a $((4 + \sqrt{2})/2 + \varepsilon)$ -approximation for every $\varepsilon > 0$.

In Section 2, we formalize our model for the movement of the points and give the definition of the diameter of a cluster in our setting, to precisely state the two variants of clustering we address. In Section 3, we present the polynomial-time algorithm for the first variant and, in Section 4, we present the approximation for the second variant and a related open problem.

2. ONE DIMENSIONAL KINETIC MODEL AND THE PROBLEMS

In our kinetic model, n points move with uniform rectilinear velocity during a continuous time interval. Without loss of generality, the time interval is $[0, 1]$. Each point $i \in \{1, 2, \dots, n\}$ has an initial position $x_i(0)$ and its velocity is given by a vector v_i . We only consider points in \mathbb{R} , so the position and the velocity are real numbers. A positive/negative velocity indicates a movement to the right/left respectively. This is a particular case of the KDS [BGH99] and Atallah's model [Ata85].

At an instant t in $[0, 1]$, the position $x_i(t)$ of a point i with initial position $x_i(0)$ and velocity v_i is given by the function

$$x_i(t) = x_i(0) + v_i t.$$

This function represents a segment on the Cartesian plane, called the *trajectory* of point i , and is given by the pair $(x_i(0), v_i)$. We draw the Cartesian plane with the horizontal axis representing the position x and the vertical axis representing the time t . Since the time interval is always $[0, 1]$, the strip of the plane between $t = 0$ and $t = 1$ will be called *time-strip*.

For our purpose, no two points have the same trajectory, or they can be treated as one. Hence, we assume a one-to-one relation between moving points and their trajectories, and mostly refer to trajectories instead of moving points in what follows.

Given a finite set S of trajectories, a *cluster* is a subset of S and a *k -clustering* is a partition of S into k clusters. Note that, as a cluster might be empty, every k' -clustering for $k' < k$ corresponds to a k -clustering by adding $k - k'$ empty clusters. Conversely, any k -clustering of S with more than $|S|$ clusters may be converted into an $|S|$ -clustering by disregarding some empty clusters. So we may assume that $1 \leq k \leq |S|$. The *left side* of a nonempty cluster C is the piecewise linear function $\min_{i \in C} x_i(t)$ for $t \in [0, 1]$.

Analogously, the *right side* is $\max_{i \in C} x_i(t)$ for $t \in [0, 1]$. The *span* of a cluster C , $\text{span}(C)$, is empty if C is empty, otherwise it is the region within the time-strip bounded by the left and right sides of C . The *diameter* of C is the area of its span, denoted by $\text{diam}(C)$. See Figure 1.




FIGURE 1. Representation of a set of moving points by their trajectories. The highlighted region is the span of the set.

The reason to consider the diameter of a cluster as the area of its span is because we deal with continuous time. Usually the diameter of the cluster is the distance between the farthest pair of points in it. So, for a continuous time interval, we integrate this distance, as shown by Equation (1), which corresponds to the area of the cluster's span:

$$\int_0^1 \left(\max_{i \in C} x_i(t) - \min_{i \in C} x_i(t) \right) dt = \text{area}(\text{span}(C)) = \text{diam}(C). \quad (1)$$

Note that one can calculate the diameter of a cluster C in time polynomial on $|C|$. If $|C| \leq 1$, then its diameter is 0. Otherwise, we can calculate the diameter of C in polynomial time by first getting a description of its span and then using any known algorithm to calculate the area of polygons. There are algorithms for this that run in time linear on the number of vertices of the span, which is linear on the number of trajectories [O'R98]. The left and right sides of C can be obtained by a divide-and-conquer algorithm that runs in $O(|C| \log |C|)$ time [Ata85].

Two problems are studied in the next sections. The *sum of diameters kinetic 1D k -clustering problem* (k -KinClust1D-SD) consists of, given a positive integer k and a finite set S of trajectories, finding a k -clustering of S whose sum of the clusters diameters is minimized. The *max diameter kinetic 1D k -clustering problem* (k -KinClust1D-MD) consists of, given a positive integer k and a finite set S of trajectories, finding a k -clustering of S whose maximum cluster diameter is minimized. The static 2-dimensional versions of these clustering problems are well-known NP-hard problems. For both, a 2-approximation is known and is best possible unless $P = NP$. The static 1D versions of both are polynomially solvable [Bru78].

3. MINIMIZING THE SUM OF THE DIAMETERS

Let \mathcal{C} be a k -clustering of a set of trajectories. We denote by

$$\text{sd}(\mathcal{C}) = \sum_{C \in \mathcal{C}} \text{diam}(C)$$

the sum of the diameters of the clusters in \mathcal{C} . With this, we can formally state our first problem as follows.

Problem 3.1 (*k-KinClust1D-SD*). Given a finite set of trajectories S and a positive integer k , find a k -clustering \mathcal{C} of S such that $\text{sd}(\mathcal{C})$ is minimum.

Let S be a finite set of n trajectories. Let \mathcal{C}^* be an optimal k -clustering of S , i.e., $\text{sd}(\mathcal{C}^*)$ is minimum among all k -clusterings of S . We denote by $\text{sd}^*(S, k) = \text{sd}(\mathcal{C}^*)$ the value of an optimal k -clustering, which is, in this case, the sum of the diameters of the clusters in \mathcal{C}^* .

For the next lemma, recall that we assumed that $k \leq n$.

Lemma 3.2. Let S be a set of n trajectories. There is always an optimal k -clustering of S for the *k-KinClust1D-SD* without empty clusters.

Proof. Let \mathcal{C}^* be an optimal k -clustering of S with minimal number of empty clusters. Suppose that \mathcal{C}^* has at least one empty cluster. Let C be a non-empty cluster of \mathcal{C}^* with at least two trajectories. Such cluster must exist as $k \leq n$.

Take some trajectory s from the left side of C . Remove s from C and replace an empty cluster by a cluster containing only s . This does not increase the diameter of C and the new cluster has diameter zero. Thus, we obtained a k -clustering \mathcal{C}' with $\text{sd}(\mathcal{C}') \leq \text{sd}(\mathcal{C}^*)$, which means $\text{sd}(\mathcal{C}') = \text{sd}(\mathcal{C}^*)$ because \mathcal{C}^* is optimal, and with fewer empty clusters, contradicting the choice of \mathcal{C}^* . Hence, there is always an optimal k -clustering of S for the *k-KinClust1D-SD* without empty clusters. \square

By Lemma 3.2, we may always look for an optimal k -clustering without empty clusters. However, k -clusterings with empty clusters can be used to bound the value of $\text{sd}^*(S, k)$.

For any integer k with $1 \leq k \leq n$, the number of distinct k -clusterings of S without empty clusters is given by the Stirling number of the second kind, denoted by $\left\{ \begin{smallmatrix} n \\ k \end{smallmatrix} \right\}$. It is known [RD69] that

$$\left\{ \begin{smallmatrix} n \\ k \end{smallmatrix} \right\} = \frac{1}{k!} \sum_{i=0}^k (-1)^{k-i} \binom{k}{i} i^n \geq \frac{1}{2} (k^2 + k + 2) k^{n-k-1} - 1 = \Omega(k^{n-k+1}).$$

For $k \geq 2$, this number is exponential on the number of trajectories even if k is a fixed value. So, it would take too long to examine every k -clustering to find an optimal one, for $k \geq 2$.

The trajectories in S divide the time-strip into convex polygonal regions. We call each of these regions a *hole* (of S). Note that a hole always has a positive area, thus its interior is never empty. The set of holes of S is denoted by $H(S)$, or simply H if S is clear from the context.

Let s be a trajectory in S and h be a hole. We say that s is *to the left* of h if there exists a point (x', t') inside h such that $x_s(t') < x'$. Otherwise, s is *to the right* of h .

Intuitively, we trace a horizontal line passing through h . Then, s is to the left of h if, in this line, the point intersecting s is to the left of some point of h . See Figure 2.




FIGURE 2. Trajectories with black ends are to the left of h and trajectories with white ends are to the right.

Note that each trajectory is either to the left or to the right of h . Thus, each hole h partitions S into two parts $S_r(h)$ and $S_\ell(h)$, with the trajectories in S to the right and to the left of h respectively. Given a k -clustering \mathcal{C} of S , a hole that is contained in the span of some cluster in \mathcal{C} is *covered* by \mathcal{C} . Otherwise it is *uncovered* by \mathcal{C} . A hole h *separates* distinct clusters C_1 and C_2 in \mathcal{C} if $C_1 \subseteq S_\ell(h)$ and $C_2 \subseteq S_r(h)$, or $C_1 \subseteq S_r(h)$ and $C_2 \subseteq S_\ell(h)$. The reference to the k -clustering is omitted when it is clear.

Since we only consider finite sets of trajectories, the number of holes in these sets is not only finite, but polynomial on the number of trajectories. This is a straightforward consequence, when we extend the trajectories to lines, of the fact that a set of n lines divides the plane in at most $n(n+1)/2 + 1$ regions. Also, for any finite set of trajectories, there are always two *unbounded* holes, h_- and h_+ , whose region extends infinitely to the left and to the right, respectively. The other holes are called *bounded*.

Lemma 3.3. Let S be a set of trajectories and let \mathcal{C}^* be an optimal k -clustering of S . For any two distinct clusters of \mathcal{C}^* , there is a hole of S separating them.

Proof. Suppose there are clusters C_1 and C_2 in \mathcal{C}^* such that there is no hole of S separating them. This means that $\text{span}(C_1) \cap \text{span}(C_2) \neq \emptyset$ and, consequently, $\text{span}(C_1) \cup \text{span}(C_2) = \text{span}(C_1 \cup C_2)$, otherwise there would be a hole of S separating C_1 and C_2 . Hence, $\text{diam}(C_1) + \text{diam}(C_2) > \text{diam}(C_1 \cup C_2)$. So, merging C_1 and C_2 would result in a better k -clustering. Therefore, any two distinct clusters in \mathcal{C}^* must be separated by a hole of S . \square

Notice that Lemma 3.3 does not guarantee that there is always an optimal k -clustering with $k - 1$ uncovered holes of S , since all holes separating two clusters could be covered by other clusters, as shown in Figure 3. However, we can guarantee the existence of at least one uncovered hole in an optimal k -clustering, for $k > 1$.

Let S be a finite set of n trajectories. We say that a trajectory s in S is *leftmost* if $x_s(0)$ is minimum in S and, in case of ties, $x_s(1)$ is minimum within the tied trajectories.

Lemma 3.4. Let S be a set of trajectories and let \mathcal{C}^* be an optimal k -clustering of S . For $k > 1$, there is at least one uncovered hole of S .




FIGURE 3. A not well-separated 3-clustering.

Proof. Suppose that every hole of S is covered by \mathcal{C}^* . Then,

$$\text{sd}(\mathcal{C}^*) = \sum_{C \in \mathcal{C}^*} \text{diam}(C) \geq \sum_{h \in H} \text{area}(h).$$

Let s be the leftmost trajectory of S . There is a hole h_s in H separating $\{s\}$ and $S \setminus \{s\}$. Consider the k -clustering \mathcal{C} with the clusters $\{s\}$, $S \setminus \{s\}$, and the remaining clusters empty. Then,

$$\text{sd}(\mathcal{C}) = \sum_{h \in H} \text{area}(h) - \text{area}(h_s) < \sum_{h \in H} \text{area}(h) \leq \text{sd}(\mathcal{C}^*).$$

This is a contradiction, since \mathcal{C}^* is an optimal k -clustering. Therefore, there is at least one hole uncovered by \mathcal{C}^* . \square

Lemmas 3.3 and 3.4 state some properties of an optimal k -clustering. So, we do not need to consider every possible k -clustering in order to find an optimal one. Considering only the k -clusterings with such properties is enough. To do that, we define a *k-good sequence* for S as a sequence $\{(h_i, C_i)\}_{i=1}^{k-1}$ such that, for $1 \leq i \leq k-1$ and $\mathcal{C}_0 = \{S\}$, we have

$$\begin{aligned} C_i &\in \mathcal{C}_{i-1} \\ h_i &\in H \text{ contained in } \text{span}(C_i) \\ \mathcal{C}_i &= (\mathcal{C}_{i-1} \setminus \{C_i\}) \cup \{(C_i)_\ell(h_i), (C_i)_r(h_i)\}. \end{aligned} \tag{2}$$

For $i \geq 1$, each \mathcal{C}_i is an $(i+1)$ -clustering obtained by separating the cluster C_i of \mathcal{C}_{i-1} using the hole h_i . Hence, a k -good sequence defines a k -clustering.

Theorem 3.5. Let S be a set of trajectories and $k \geq 1$. Every optimal k -clustering of S is defined by a k -good sequence.

Proof. The proof is by induction on k . For $k = 1$, the statement is trivial, since $\mathcal{C}_0 = \{S\}$ is the only 1-clustering of S . So, consider $k > 1$ and let \mathcal{C}^* be an optimal k -clustering of S .

By Lemma 3.4, there is an uncovered hole h_1 in H . Then, h_1 separates S into $S_1 = S_\ell(h_1)$ and $\bar{S}_1 = S_r(h_1)$, so that $C \subseteq S_1$ or $C \subseteq \bar{S}_1$, for every $C \in \mathcal{C}^*$. Let $\mathcal{D} = \{C \in \mathcal{C}^* \mid C \subseteq S_1\}$, $k_1 = |\mathcal{D}|$, $\bar{\mathcal{D}} = \{C \in \mathcal{C}^* \mid C \subseteq \bar{S}_1\}$, and $\bar{k}_1 = |\bar{\mathcal{D}}|$. Note that $k_1 \geq 1$, $\bar{k}_1 \geq 1$, and $k_1 + \bar{k}_1 = k$. Moreover, \mathcal{D} is an optimal k_1 -clustering of S_1 and $\bar{\mathcal{D}}$ is an optimal \bar{k}_1 -clustering of \bar{S}_1 , otherwise the k -clustering $\mathcal{C}^* = \mathcal{D} \cup \bar{\mathcal{D}}$ would not be optimal for S .

By induction hypothesis, there is a k_1 -good sequence $\{(d_i, C_i)\}_{i=1}^{k_1-1}$ defining \mathcal{D} and there is a \bar{k}_1 -good sequence $\{(\bar{d}_i, \bar{C}_i)\}_{i=1}^{\bar{k}_1-1}$ defining $\bar{\mathcal{D}}$. Note that d_i is a hole of S_1 . If d_i is not also a hole of S , then d_i corresponds to the union of two or more holes of S . Any one of these holes of S inside d_i separates S_1 in the same way as d_i . Thus, we can exchange d_i by one of the holes of S inside it. The same can be done to each hole \bar{d}_i of \bar{S}_1 .

Therefore, $(h_1, S), (d_1, C_1), \dots, (d_{k_1-1}, C_{k_1-1}), (\bar{d}_1, \bar{C}_1), \dots, (\bar{d}_{\bar{k}_1-1}, \bar{C}_{\bar{k}_1-1})$ is a k -good sequence that defines \mathcal{C}^* . \square

Lemma 3.6. Let S be a set of n trajectories and $k \geq 1$. The number of k -good sequences of S is $O(n^{2(k-1)}(k-1)!)$.

Proof. Remember that $|H| = O(n^2)$. For each element (h_i, C_i) of a k -good sequence, there are at most $|H| - i + 1$ choices for hole h_i , and at most $|C_{i-1}| = i$ choices for cluster C_i . Thus, the number of such sequences is at most

$$\prod_{i=1}^{k-1} (|H| - i + 1) i = O\left(\prod_{i=1}^{k-1} n^2 i\right) = O(n^{2(k-1)}(k-1)!). \quad \square$$

So, instead of considering all the $\binom{n}{k} = \Omega(k^{n-k})$ possible k -clusterings of S , we may consider only those that are defined by a k -good sequence. Lemma 3.6 guarantees that the number of such sequences is polynomial in the input size for a fixed value of k . Moreover, since the diameter of a cluster with m trajectories can be calculated in time $O(m \log m)$, an algorithm that examines every k -clustering defined by a k -good sequence runs in time $O((k-1)! n^{2k-1} \log n)$.

3.1. Well-separated clusterings. Depending on the application, we may want the similarity between distinct clusters to be as small as possible. For example, a cluster whose span is contained in the span of another one may be undesirable even if this configuration was the only one to minimize the objective function. We say that a k -clustering is *well-separated* if its clusters are pairwise separated by an uncovered hole. Figure 3 shows an example of a 3-clustering that is not well-separated, because there is no uncovered hole separating the bold trajectory from the light gray cluster.

If we require the holes b_i of a k -good sequence to be all uncovered, then the algorithm described in Lemma 3.6 would find an optimal well-separated k -clustering for k -KinClust1D-SD. However, we can do better using a dynamic programming algorithm that is polynomial in the input size, even if k is part of the input.

Let $\mathcal{C}_H = \cup_{h \in H} \{S_\ell(h), S_r(h)\}$. Consider the partial order \preceq over \mathcal{C}_H defined as follows: for every C_1 and C_2 in \mathcal{C}_H , we say $C_1 \preceq C_2$ if and only if $C_1 \subseteq C_2$. We denote by D_H the directed acyclic graph (dag) representing this partial order. See Figure 4. Notice that $|V(D_H)| = |\mathcal{C}_H| \leq 2|H|$.

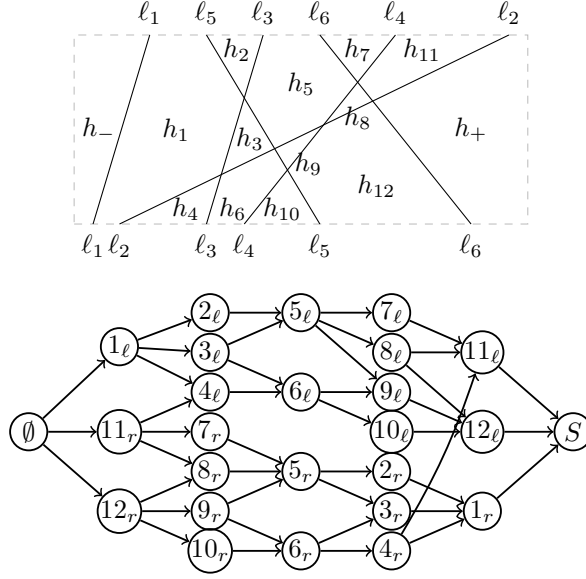


FIGURE 4. Dag D_H . Vertices i_ℓ and i_r represent $S_\ell(h_i)$ and $S_r(h_i)$, respectively.

Lemma 3.7. Let S be a set of trajectories and \mathcal{C} be a well-separated k -clustering of S . There is an ordering C_1, C_2, \dots, C_k of the clusters in \mathcal{C} such that, for every $1 \leq i \leq k-1$,

$$\bigcup_{j=1}^i C_j \quad \text{and} \quad \bigcup_{j=i+1}^k C_j$$

are separated by some uncovered hole h_i of H .

Proof. Let \tilde{H} be the set of holes of S uncovered by \mathcal{C} . Suppose that $k \geq 2$, otherwise the statement is trivial. Consider the following construction. Let

$$A_1 = \arg \min_{A \in \mathcal{C}_{\tilde{H}}} |A|, \quad \text{and} \quad A_i = \arg \min_{\substack{A \in \mathcal{C}_{\tilde{H}}, \\ A_{i-1} \subset A}} |A|, \quad \text{for } 2 \leq i \leq k-1.$$

Also, consider a hole $h_i \in \tilde{H}$ such that $A_i \in \{S_\ell(h_i), S_r(h_i)\}$, for $1 \leq i \leq k-1$.

Take $C_1 = A_1$, $C_i = A_i \setminus A_{i-1}$, for each $2 \leq i \leq k-1$, and $C_k = S \setminus A_{k-1}$. Observe that $\bigcup_{j=1}^i C_j$ and $\bigcup_{j=i+1}^k C_j$ are separated by h_i , for all $1 \leq i \leq k-1$. Now, we only need to show that this construction is well-defined, i.e., none of the sets is empty and $\mathcal{C} = \{C_1, C_2, \dots, C_k\}$.

Since \mathcal{C} is well-separated, $|\tilde{H}| \geq k-1 \geq 1$. Hence, C_1 exists and is nonempty. Moreover, C_1 is a cluster of \mathcal{C} , otherwise there would be at least two clusters of \mathcal{C} in C_1 . However, these two clusters would have to be separated by an uncovered hole h , implying that $\min(|S_\ell(h)|, |S_r(h)|) < |C_1|$. This contradicts the choice of C_1 .

Fix $2 \leq i \leq k-1$. Suppose that C_1, C_2, \dots, C_{i-1} are nonempty and that they are clusters of \mathcal{C} . Since \mathcal{C} is a well-separated k -clustering, the other $k-i+1 \geq 2$ clusters of \mathcal{C} are inside $S \setminus A_{i-1}$. Hence, there is a hole h in \tilde{H} separating some pair of clusters inside $S \setminus A_{i-1}$. Thus, we can conclude that A_{i-1} is properly contained in $S_\ell(h)$ or in $S_r(h)$, implying that A_i and C_i are nonempty. Since A_i has minimum cardinality and $A_{i-1} \subset A_i$,

by the same argument used previously for C_1 , the set $A_i \setminus A_{i-1}$ indeed is a cluster of \mathcal{C} . Therefore, $\mathcal{C} = \{C_1, C_2, \dots, C_k\}$. \square

A k -chain is a sequence $\{C_i\}_{i=1}^{k-1}$ of $k-1$ distinct elements of $\mathcal{C}_H \setminus \{S, \emptyset\}$ such that $C_1 \preceq C_2 \preceq \dots \preceq C_{k-1}$. Note that, for any k -chain, $C_{k-1} \preceq S$. A k -chain defines the k -clustering $\{C_1, C_2 \setminus C_1, C_3 \setminus C_2, \dots, C_{k-1} \setminus C_{k-2}, S \setminus C_{k-1}\}$.

Theorem 3.8. Let S be a set of trajectories. Every well-separated k -clustering of S is defined by a k -chain.

Proof. Let $\mathcal{C} = \{C_1, C_2, \dots, C_k\}$ be a well-separated k -clustering of S . Suppose that the clusters in \mathcal{C} are sorted as in Lemma 3.7.

Thus, for every $1 \leq i \leq k-1$, we have that $S_i = \cup_{j=1}^i C_j \in \{S_\ell(h_i), S_r(h_i)\}$ for some uncovered hole h_i . Consider $S_k = S$. Hence, $S_i \in \mathcal{C}_H$ and $S_i \subseteq S_{i+1}$ for every $1 \leq i \leq k-1$. Therefore, $\{\cup_{j=1}^i C_j\}_{i=1}^{k-1}$ is a k -chain that defines \mathcal{C} . \square

Notice that \emptyset and S are the source and sink of D_H , respectively. For each C in \mathcal{C}_H and each $j \geq 1$, let $\text{sd-rec}(C, j)$ be the sum of the clusters diameters of an optimal well-separated j -clustering of $S \setminus C$ for $k\text{-KinClust1D-SD}$. The following recurrence holds.

$$\text{sd-rec}(C, j) = \begin{cases} 0, & \text{if } C = S \\ \text{diam}(S \setminus C), & \text{if } j = 1 \\ \min_{C' \preceq C} \{\text{diam}(C' \setminus C) + \text{sd-rec}(C', j-1)\}, & \text{otherwise.} \end{cases} \quad (3)$$

Given a set of n trajectories S and an integer k , with $1 \leq k \leq n$, the minimum sum of the diameters of an optimal well-separated k -clustering of S is $\text{sd-rec}(\emptyset, k)$. A straightforward dynamic programming implementation of recurrence (3) to calculate $\text{sd-rec}(\emptyset, k)$ consists in filling a $|\mathcal{C}_H| \times k$ matrix, whose rows represent elements of \mathcal{C}_H and columns represent values of j in recurrence (3). To fill in each matrix position from row S takes time $O(1)$, and from column $j = 1$ takes time $O(n \log n)$, which is the time to calculate the diameter of a set of trajectories. The remaining matrix positions are filled in an order such that, when filling in position (C, j) , the values of $\text{sd-rec}(C', j')$, for all $C' \preceq C$ and $j' \leq j$, are already filled. So it takes time $O(|\mathcal{C}_H| n \log n)$ to calculate each of these positions. Therefore, it is possible to calculate $\text{sd-rec}(\emptyset, k)$ in time $O(k|\mathcal{C}_H|^2 n \log n) = O(n^6 \log n)$. After the whole matrix is filled, we can construct an optimal well-separated k -clustering of S in time $O(k|\mathcal{C}_H| n \log n)$ by tracing back recurrence (3).

3.2. Approximation using well-separated k -clusterings. Let S be a set of trajectories and \mathcal{C}^* be an optimal k -clustering of S . We say that \mathcal{C}' is an optimal well-separated k -clustering if $\text{sd}(\mathcal{C}')$ is minimal among all well-separated k -clusterings of S . As shown in Figure 3, there are cases in which no optimal k -clustering for $k\text{-KinClust1D-SD}$ is well-separated.

Next, we prove a bound on the value of an optimal well-separated k -clustering in terms of the value of an optimal k -clustering. This bound allows one to conclude that the algorithm presented in Section 3.1 is an $O(k)$ -approximation for the k -KinClust1D-SD.

Theorem 3.9. Let S be a set of $n \geq 3$ trajectories and \mathcal{C}^* be an optimal k -clustering of S . If \mathcal{C}' is an optimal well-separated k -clustering of S for the k -KinClust1D-SD, then

$$\text{sd}(\mathcal{C}') \leq (1 + \lfloor k/2 \rfloor) \text{sd}(\mathcal{C}^*).$$

Proof. For $k = 1$, clearly $\text{sd}(\mathcal{C}^*) = \text{sd}(\mathcal{C}')$. By Lemma 3.4, this equality is also true for $k = 2$. Thus, we may assume that $k \geq 3$.

By Lemma 3.4, there is at least one hole of S uncovered by \mathcal{C}^* . The holes uncovered by \mathcal{C}^* partition S into $\mathcal{D} = \{D_1, D_2, \dots, D_\ell\}$ with $\ell < k$. Notice that \mathcal{D} is well-separated. Moreover, each cluster in \mathcal{C}^* is contained in some D in \mathcal{D} .

We denote by \mathcal{C}_D^* the subset of clusters $\{C \in \mathcal{C}^* \mid C \subseteq D\}$. For each $1 \leq j \leq \ell$, let H_j be the set of holes of S that separate a pair of clusters in $\mathcal{C}_{D_j}^*$. Each hole in H_j is covered by \mathcal{C}^* , since it separates clusters contained in the same D of \mathcal{D} .

For each D_j in \mathcal{D} ,

$$\text{diam}(D_j) \leq \sum_{C \in \mathcal{C}_{D_j}^*} \text{diam}(C) + \sum_{h \in H_j} \text{area}(h).$$

Hence,

$$\text{sd}(\mathcal{D}) = \sum_{j=1}^{\ell} \text{diam}(D_j) \leq \sum_{C \in \mathcal{C}^*} \text{diam}(C) + \sum_{j=1}^{\ell} \sum_{h \in H_j} \text{area}(h) \leq \text{sd}(\mathcal{C}^*) + \frac{k}{2} \text{sd}(\mathcal{C}^*).$$

The last inequality comes from the fact that, if a hole h belongs to some H_j , then $|\mathcal{C}_{D_j}^*| \geq 2$, since h separates two clusters in $\mathcal{C}_{D_j}^*$. So, each hole of S can appear in at most $\lfloor k/2 \rfloor$ different sets H_j .

Therefore, $\text{sd}(\mathcal{C}') \leq \text{sd}(\mathcal{D}) \leq (1 + \lfloor k/2 \rfloor) \text{sd}(\mathcal{C}^*)$. \square

Notice that, if the sets H_j in the proof of Theorem 3.9 were pairwise disjoint, each hole in $\bigcup_{j=1}^{\ell} H_j$ would appear only once in the summation, implying that

$$\sum_{j=1}^{\ell} \sum_{h \in H_j} \text{area}(h) \leq \text{sd}(\mathcal{C}^*).$$

Thus, we would have $\text{sd}(\mathcal{C}') \leq 2 \text{sd}(\mathcal{C}^*)$.

4. MINIMIZING THE MAXIMUM DIAMETER

Let \mathcal{C} be a k -clustering of a set of trajectories. We denote by

$$\text{md}(\mathcal{C}) = \max_{C \in \mathcal{C}} \text{diam}(C)$$

the maximum diameter of \mathcal{C} . The second k -clustering problem that we consider is the following.

Problem 4.1 (*k-KinClust1D-MD*). Given a finite set of trajectories S and a positive integer k , find a k -clustering \mathcal{C} of S such that $\text{md}(\mathcal{C})$ is minimum.

Let S be a finite set of trajectories. Let \mathcal{C}^* be an optimal k -clustering of S , i.e., $\text{md}(\mathcal{C}^*)$ is minimal among all k -clusterings of S . We denote by $\text{md}^*(S, k) = \text{md}(\mathcal{C}^*)$ the value of an optimal k -clustering, which is, in this case, the maximum diameter of \mathcal{C}^* .

Unlike the *k-KinClust1D-SD*, there are cases in which no optimal k -clustering for *k-KinClust1D-MD* is defined by a k -good sequence. Figure 5 shows an example for $k = 2$. The 2-clustering of the example has no uncovered hole. Thus, this clustering cannot be defined by a k -good sequence. Both clusters have diameter 1 and it is not difficult to check that any cluster with three trajectories has diameter greater than 1. Also, any other 2-clustering has maximum diameter greater than 1. Hence, the 2-clustering of the example is optimal for the *k-KinClust1D-MD*. Therefore, we cannot use the same approach used for the *k-KinClust1D-SD*, for a constant value of k .




FIGURE 5. Different end types represent different clusters. This is the only optimal 2-clustering for *k-KinClust1D-MD* and it is not defined by any k -good sequence.

However, if we are specifically looking for an optimal well-separated k -clustering for *k-KinClust1D-MD*, a simple adaptation of recurrence (3) can be used, resulting in an algorithm that finds an optimal well-separated k -clustering for *k-KinClust1D-MD* in time $O(k|\mathcal{C}_H|^2 n \log n)$. For the *k-KinClust1D-MD* without restriction on the clustering, we show some approximation results in the next subsections.

4.1. First approximation. Our first result for the *k-KinClust1D-MD* is achieved by a reduction to the classical (metric) k -center problem.

Problem 4.2 (*kCenter*). Given a positive integer k , a finite set S , and a distance function d over S , find a subset X of S such that $|X| \leq k$ and $\max_{s \in S} \min_{x \in X} d(s, x)$ is minimum.

The *kCenter* is known to be NP-hard [Gon85, KH79], but there are simple 2-approximations [Gon85, HS86]. Moreover, no better approximation factor is possible unless $P = NP$ [Gon85, HN79].

The elements of the desired set X are called *centers*. We basically want to find at most k centers such that the largest distance between an element of S and its nearest center is minimized. Note that, for each center x in X , we have an *induced* cluster

$$C_x = \{s \in S \mid d(s, x) \leq d(s, x'), \text{ for all } x' \in X\}.$$

A k -clustering induced by X is roughly $\mathcal{C}_X = \{C_x \mid x \in X\}$. In fact, at first, \mathcal{C}_X may not be a k -clustering since its clusters are not necessarily disjoint. If this is the case, that is, if $C_{x_1} \cap C_{x_2} \neq \emptyset$ for some x_1 and x_2 in X , then we can just remove this intersection from either one of the clusters.

From an instance (S, k) of the k -KinClust1D-MD, we build an instance of the k Center by setting the distance between two trajectories s_1 and s_2 in S as $\text{diam}(\{s_1, s_2\})$. These distances satisfy the triangle inequality, as shown by Lemma 4.3, so the instance is metric and we can apply any 2-approximation for the k Center, outputting a k -clustering induced by the k selected centers.

Lemma 4.3. Let S be a set of trajectories. For any three trajectories a , b , and c of S , we have $\text{diam}(\{a, b\}) + \text{diam}(\{b, c\}) \geq \text{diam}(\{a, c\})$.

Proof. Given two trajectories of S , say r and s , we have that

$$\text{diam}(\{r, s\}) = \int_0^1 \left(\max_{i \in \{r, s\}} x_i(t) - \min_{i \in \{r, s\}} x_i(t) \right) dt = \int_0^1 |x_r(t) - x_s(t)| dt.$$

By the linearity of integrals and the subadditivity of modulus, we have that

$$\begin{aligned} \text{diam}(\{a, b\}) + \text{diam}(\{b, c\}) &= \int_0^1 (|x_a(t) - x_b(t)| + |x_b(t) - x_c(t)|) dt \\ &\geq \int_0^1 |x_a(t) - x_c(t)| dt \\ &= \text{diam}(\{a, c\}). \end{aligned} \quad \square$$

Let s be a trajectory in S . To facilitate the reading, in what follows we abuse notation and consider $s^t = x_s(t)$ as the position of s at time t and also as a point $s^t = (x_s(t), t)$ in the plane. In particular, for $t = 1/2$ we write \bar{s} instead of $s^{1/2}$. The precise meaning of s^t will be clear from the context.

Lemma 4.4. Let s and v be two trajectories. If $|\bar{s} - \bar{v}| > r$ for some $r > 0$, then $\text{diam}(\{s, v\}) > r$.

Proof. Let p be a trajectory parallel to s that passes through \bar{v} . Hence, $\text{diam}(\{s, p\}) = |s^t - p^t| > r$, for any $t \in [0, 1]$. If $v = p$, we are done. So, suppose that $v \neq p$.

Since p and v intersect at $t = 1/2$, we have that $|v^0 - p^0| = |v^1 - p^1|$. Thus, $\text{span}(\{v, p\})$ consists of two congruent triangles with vertex \bar{v} in common. Assume, without loss of generality, that $\bar{s} < \bar{v}$ and $v^0 > p^0$, as shown in Figure 6.

Let A , B , C , and D be the regions indicated in Figure 6. If s does not intersect v or if s intersects v at one of its extremities, then A is empty. Note that

$$\begin{aligned} \text{diam}(\{s, v\}) &= \text{area}(A) + \text{area}(C) + \text{area}(D), \\ \text{diam}(\{p, v\}) &= \text{area}(A) + \text{area}(B) + \text{area}(D), \\ \text{diam}(\{s, p\}) &= \text{area}(B) + \text{area}(C). \end{aligned}$$




FIGURE 6. Calculating $\text{diam}(\{s, v\})$ as a function of $\text{diam}(\{p, v\})$.

Moreover, as the two triangles in the $\text{span}(\{v, p\})$ are congruent, we have $\text{area}(D) = \text{area}(A) + \text{area}(B)$ and, consequently, $\text{area}(D) \geq \text{area}(B)$. Therefore

$$\begin{aligned} \text{diam}(\{s, v\}) &= \text{area}(A) + (\text{diam}(\{s, p\}) - \text{area}(B)) + \text{area}(D) \\ &\geq \text{diam}(\{s, p\}) > r. \end{aligned} \quad \square$$

Lemma 4.5. Let S be a set of trajectories and s be a trajectory of S . If r is such that $\text{diam}(\{s, a\}) \leq r$ for all a in S , then $\text{diam}(S) \leq (2 + \sqrt{2})r$.

Proof. Notice that, for any trajectory a such that $|a^0 - s^0| > (1 + \sqrt{2})r$, we have $\text{diam}(\{s, a\}) > r$, so a does not belong to S . Any trajectory a in S such that $|a^0 - s^0| > 2r$ intersects s and $|a^1 - s^1| < 2r$. Both observations are true if we exchange $a^0 - s^0$ with $a^1 - s^1$.

By Lemma 4.4, for every trajectory a in S , we have $|\bar{s} - \bar{a}| \leq r$. Also, the left side of S is convex and the right side of S is concave with respect to the vertical axis. Hence, we can conclude that the span of S is inside the polygon P highlighted in Figure 7.




FIGURE 7. Polygon P containing the span of S .

The area of P is given by the sum of the area of two symmetric trapezoids of height $1/2$:

$$\text{area}(P) = 2 \left(\frac{1}{2} \cdot \frac{2(1 + \sqrt{2})r + 2r}{2} \right) = (2 + \sqrt{2})r.$$

Therefore $\text{diam}(S) \leq \text{area}(P) \leq (2 + \sqrt{2})r$. \square

If the maximum distance to a center chosen by a 2-approximation for the k Center is r , then $r \leq 2 \text{md}^*(S, k)$. Indeed, if a cluster C is such that $\text{diam}(C) \leq q$, then any two trajectories in C are at a distance at most q . Since, in a cluster produced by the algorithm, any two trajectories are at a distance at most r from the center, then $q \leq 2r$.

Using the lemma and the bound on r , we deduce that the described algorithm is a $2(2 + \sqrt{2})$ -approximation.

4.2. Second approximation. Next we describe a better approximation for k -*KinClust1D-MD*, inspired in the bottleneck method of Hochbaum and Shmoys [HS86]. First we describe a greedy algorithm that, given a set S of n trajectories and a positive number $D \geq 0$, obtains a k' -clustering \mathcal{C}_D with $\text{md}(\mathcal{C}_D) \leq 2.71 D$. If $k' > k$, we will have a certificate that $\text{md}^*(S, k) > D$. If $k' = k$, the algorithm gives a 2.71-approximation.

Remember that a trajectory s in S is the leftmost if $x_s(0)$ is minimum in S and, in case of ties, $x_s(1)$ is minimum within the tied trajectories.

```

GREEDYPARTITION( $S, D$ )
  if  $S = \emptyset$  then return  $\emptyset$ 
  let  $s$  be the leftmost trajectory in  $S$ 
   $C \leftarrow \{s\}$ 
  for each  $s' \in S$  do
    if  $\text{diam}(\{s, s'\}) \leq D$  then  $C \leftarrow C \cup \{s'\}$ 
  return  $\{C\} \cup \text{GREEDYPARTITION}(S \setminus C, D)$ 

```

It is clear that the algorithm GREEDYPARTITION terminates. Since at each call at least one trajectory is removed from the set of trajectories, the algorithm does at most $n - 1$ recursive calls. Each call takes time $O(n)$ excluding the recursive call, since the diameter of a subset of two trajectories can be calculated in constant time. Finding the leftmost trajectory in S can also be done in constant time if we preprocess S in time $O(n \log n)$, sorting the trajectories accordingly. Thus, GREEDYPARTITION has time complexity $O(n^2)$.

Lemma 4.6. Each cluster built by GREEDYPARTITION(S, D) has diameter at most $(4 + \sqrt{2})D/2$.

Proof. The proof of this lemma is similar to the proof of Lemma 4.5, but here we have a better bound on the span of the cluster. Let s be the leftmost trajectory in S and C be the cluster returned by GREEDYPARTITION(S, D) that contains s . Any trajectory u in C has $u^0 \geq s^0$, since s is the leftmost one. Also, we have that $\text{diam}(\{s, u\}) \leq D$, so $s^1 - (1 + \sqrt{2})D \leq u^1 \leq s^1 + 2D$ and $u^0 \leq s^0 + (1 + \sqrt{2})D$. Thus, the span of C is contained in the polygon P highlighted in Figure 8, whose area is the sum of two trapezoids of height $1/2$:

$$\text{area}(P) \leq \frac{((1 + \sqrt{2})D + 2D) \frac{1}{2}}{2} + \frac{((2D + (1 + \sqrt{2})D) + 2D) \frac{1}{2}}{2} = \frac{(4 + \sqrt{2})D}{2}.$$

Therefore, $\text{diam}(C) \leq \text{area}(P) \leq (4 + \sqrt{2})D/2$. \square

Lemma 4.7. Let S be a set of trajectories. Let k be a positive integer and D a positive real. If GREEDYPARTITION(S, k) returns a k' -clustering with $k' > k$, then $\text{md}^*(S, k) > D$.




FIGURE 8. Polygon containing the span of C .

Proof. Let $\mathcal{C} = \{C_1, C_2, \dots, C_{k'}\}$ be the clustering returned by $\text{GREEDYPARTITION}(S, D)$. For $1 \leq i \leq k'$, let s_i be the leftmost trajectory in C_i . If $k' > k$, then $\{s_1, s_2, \dots, s_{k'}\}$ is a certificate that $\text{md}^*(S, k) > D$.

Let s_i and s_j be any two distinct trajectories in $\{s_1, s_2, \dots, s_{k'}\}$. Since s_i and s_j belong to different clusters, $\text{diam}(\{s_i, s_j\}) > D$, otherwise GREEDYPARTITION would have put both in the same cluster. Therefore, any clustering \mathcal{C}' of S with $\text{md}(\mathcal{C}') \leq D$ must have at least $k' > k$ clusters. \square

As shown by Lemmas 4.6 and 4.7, GREEDYPARTITION is similar to what Hochbaum and Shmoys [HS86] call a relaxed decision procedure for bottleneck problems. However, we do not know how to restrict the possible values of $\text{md}^*(S, k)$ to a set of polynomial size. Thus we do an approximate binary search to get the parameter D as close as we want to $\text{md}^*(S, k)$. This is done by the algorithm $k\text{-CLUSTERINGBS}_\varepsilon$ below, for any $\varepsilon > 0$.

$k\text{-CLUSTERINGBS}_\varepsilon(S, k)$

$a \leftarrow 0$

$b \leftarrow \text{diam}(S)$

$\delta \leftarrow \frac{2\varepsilon}{4+\sqrt{2}} \min_{u,v \in S, u \neq v} \text{diam}(\{u, v\})$

while $(b - a > \delta)$ *do*

$D \leftarrow \frac{a+b}{2}$

$\mathcal{C} \leftarrow \text{GREEDYPARTITION}(S, D)$

if $|\mathcal{C}| > k$ *then* $a \leftarrow D$ *else* $b \leftarrow D$

if $|\mathcal{C}| > k$ *then return* $\text{GREEDYPARTITION}(S, b)$ *else return* \mathcal{C}

After i iterations, the search interval is halved i times and, thus, has length $2^{-i} \text{diam}(S)$. So, it takes at most $\log \text{diam}(S) - \log \delta$ iterations, for δ as in the algorithm, to decrease the search interval length to δ .

Theorem 4.8. Let S be a set of n trajectories and k be a positive integer. For every $\varepsilon > 0$, $k\text{-CLUSTERINGBS}_\varepsilon(S, k)$ is a $((4 + \sqrt{2})/2 + \varepsilon)$ -approximation for $k\text{-KinClust1D-MD}$.

Proof. First, notice that the algorithm always returns a k -clustering, since the value of b in the algorithm starts with $\text{diam}(S)$. Also, from the binary search, we know that $\text{md}^*(S, k)$ is in the interval $[a, b]$ and at the end of the search we have $b - a \leq \delta$. Thus, $b \leq \text{md}^*(S, k) + \delta$.

Suppose that $k < n$, otherwise it is trivial to find an optimal solution. Hence, one of the clusters in any optimal k -clustering must have at least two trajectories and $\text{md}^*(S, k) \geq \min_{u,v \in S, u \neq v} \text{diam}(\{u, v\})$.

Since $\delta = (2\varepsilon/(4 + \sqrt{2})) \min_{u,v \in S, u \neq v} \text{diam}(\{u, v\})$,

$$\begin{aligned} \text{md}(\mathcal{C}) &\leq \frac{4 + \sqrt{2}}{2} b \leq \frac{4 + \sqrt{2}}{2} (\text{md}^*(S, k) + \delta) \\ &= \frac{4 + \sqrt{2}}{2} \text{md}^*(S, k) + \frac{4 + \sqrt{2}}{2} \frac{2\varepsilon}{4 + \sqrt{2}} \min_{\substack{u,v \in S \\ u \neq v}} \text{diam}(\{u, v\}) \\ &= \frac{4 + \sqrt{2}}{2} \text{md}^*(S, k) + \varepsilon \min_{\substack{u,v \in S \\ u \neq v}} \text{diam}(\{u, v\}) \\ &\leq \left(\frac{4 + \sqrt{2}}{2} + \varepsilon \right) \text{md}^*(S, k). \end{aligned}$$

By the choice of δ , $k\text{-CLUSTERINGBS}_\varepsilon(S, k)$ terminates within at most $\log \text{diam}(S) + \log(4 + \sqrt{2}) - \log(2\varepsilon \min_{u,v \in S, u \neq v} \text{diam}(\{u, v\}))$ iterations. Each iteration consists of a call to `GREEDYPARTITION` which takes time $O(n^2)$. Thus, for a fixed ε , the algorithm runs in time polynomial in the input size.

Therefore, $k\text{-CLUSTERINGBS}_\varepsilon(S, k)$ is a $((4 + \sqrt{2})/2 + \varepsilon)$ -approximation for $k\text{-KinClust1D-MD}$, for every $\varepsilon > 0$. \square

Notice that the value of ε is not considered as part of the input when we say that $k\text{-CLUSTERINGBS}_\varepsilon$ runs in time polynomial in the input size. Fortunately, in the case of $k\text{-CLUSTERINGBS}_\varepsilon$, the time complexity is not only polynomial in the input size but also in $1/\varepsilon$.

5. FINAL COMMENTS

The formalization and study of kinetic versions of clustering problems seems quite intriguing and challenging. Of course it would be nice to address variants of the problems we addressed in higher dimensions, considering points moving in the plane, or in a 3D space. Also, allowing the clusters to change with time, in a smooth way, seems reasonable and leads to interesting questions.

We focused on the one-dimensional case, presenting polynomial-time algorithms for the $k\text{-KinClust1D-SD}$ under some assumptions, and approximation algorithms for the $k\text{-KinClust1D-MD}$. However, the complexity of the $k\text{-KinClust1D-SD}$ and the $k\text{-KinClust1D-MD}$ remains open, that is, we do not know whether the $k\text{-KinClust1D-SD}$ and the $k\text{-KinClust1D-MD}$ are NP-hard. It would be nice to settle the complexity of these two problems, either by proving that they are NP-hard, or by presenting polynomial-time algorithms to solve them. Meanwhile, achieving better approximations for $k\text{-KinClust1D-MD}$ and for $k\text{-KinClust1D-SD}$ would also be nice. In particular, as far as we know, it is possible that the algorithm described at the end of Section 3, that outputs an optimal

well-separated k -clustering, achieves a constant approximation ratio for k -KinClust1D-SD, that is, independent of the value of k .

REFERENCES

- [Ata85] M. Atallah, *Some dynamic computational geometry problems*, Comput. Math. Appl. **11** (1985), no. 12, 1171–1181.
- [BGH99] J. Basch, L. Guibas, and J. Hershberger, *Data Structures for Mobile Data*, J. Algorithms **31** (1999), 1–28.
- [Bru78] P. Brucker, *On the Complexity of Clustering Problems*, Optimization and Operations Research, LNEMS, vol. 157, Springer Berlin Heidelberg, 1978, pp. 45–54 (English).
- [GGH⁺03] J. Gao, L. Guibas, J. Hershberger, L. Zhang, and A. Zhu, *Discrete Mobile Centers*, Discrete Comput. Geom. **30** (2003), no. 1, 45–63 (English).
- [Gon85] T. Gonzalez, *Clustering to minimize the maximum intercluster distance*, Theor. Comput. Sci. **38** (1985), 293–306.
- [HN79] W.L. Hsu and G. Nemhauser, *Easy and hard bottleneck location problems*, Discrete Appl. Math. **1** (1979), no. 3, 209–215.
- [HP04] S. Har-Peled, *Clustering Motion*, Disc. Comput. Geom. **31** (2004), no. 4, 545–565 (English).
- [HS86] D. Hochbaum and D. Shmoys, *A Unified Approach to Approximation Algorithms for Bottleneck Problems*, J. ACM **33** (1986), no. 3, 533–550.
- [KH79] O. Kariv and S. Hakimi, *An Algorithmic Approach to Network Location Problems. I: The p -Centers*, SIAM J. Appl. Math. **37** (1979), no. 3, 513–538.
- [LHW07] J.G. Lee, J. Han, and K.Y. Whang, *Trajectory Clustering: A Partition-and-group Framework*, SIGMOD’07, ACM, 2007, pp. 593–604.
- [O’R98] J. O’Rourke, *Computational Geometry in C*, Cambridge Tracts in Theoretical Computer Science, Cambridge University Press, 1998.
- [RD69] Basil Cameron Rennie and Annette Jane Dobson, *On Stirling Numbers of the Second Kind*, Journal of Combinatorial Theory **7** (1969), no. 2, 116–121.